

How to Be a Good Graduate Student Advisor

Marie desJardins
marie@erg.sri.com

March 1994

Abstract

This paper attempts to raise some issues that are important for graduate students to be successful and to get as much out of the process as possible, and for advisors who wish to help their students be successful. The intent is not to provide prescriptive advice – no formulas for finishing a thesis or twelve-step programs for becoming a better advisor are given – but to raise awareness on both sides of the advisor-student relationship as to what the expectations are and should be for this relationship, what a graduate student should expect to accomplish, common problems, and where to go if the advisor is not forthcoming.

1 Introduction

This article originated with a discussion I had with several women professors about the problems women face in graduate school, and how more women could be encouraged to go to graduate school in computer science. Eventually, the conversation turned to the question of what these women could do in their interactions with women students to support and encourage them. I volunteered that over the course of my graduate career I had collected a variety of papers and e-mail discussions about how to be a good advisor, how to get through graduate school, and issues facing women. They were eager to get this material, and I told them I would sort through it when I got a chance.

After mentioning this project to a number of people, both graduate students and faculty –

all of whom expressed an interest in anything I could give them – I realized two things: first, the issues that we were talking about really were not just women’s issues but were of interest to all graduate students, and to all caring advisors. Second, in order to disseminate the information I had collected (and was starting to collect from others) it seemed to make more sense to compile a bibliography, and write a paper that would summarize the most useful advice and suggestions I had collected.

I solicited inputs from friends and colleagues via mailing lists and Internet bulletin boards, and collected almost an overwhelming amount of information. Sorting through it and attempting to distill the collective wisdom of dozens of articles and hundreds of e-mail messages has not been an easy task, but I hope that the results provide a useful resource for graduate students and advisors alike. The advice I give here is directed towards Ph.D. students in computer science and their advisors, since that is my background, but I believe that much of it applies to graduate students in other areas as well.

In my experience, the two main things that make graduate school hard are the unstructured nature of the process, and the lack of information about what you should spend your time on. I hope that this article will provide information for both graduate students and advisors that will help make the process less painful.

I owe a debt of gratitude to David Chapman, whose paper ([Chapman, 1988]) was an invaluable reference for me not only during the writing of this article, but during graduate school as well.

The goals of this article are to raise awareness

of the need for a healthy and interactive graduate student-advisor relationship, to provide pointers and guidance for both advisors and graduate students in navigating the maze of a doctoral degree, and to give references and resources for those who hope to learn more.

2 Before You Start

Many headaches can be avoided by doing some advance planning. First, why go to graduate school at all? The usual reasons given are that a Ph.D. is required or preferred for some jobs, especially research and academic positions; that it gives you a chance to learn a great deal about a specific area; and that it provides an opportunity to develop ideas and perform original research. Wanting to delay your job hunt is probably not a good enough reason. Graduate school is a lot of work and requires strong motivation and focus. You have to really want to be there to make it through.

It helps to have a good idea of what area you want to specialize in, and preferably a couple of particular research projects you might like to work on. Look for books and current journals and conference proceedings in your area, and read through them to get an idea of who's doing what where. (You'll be doing a *lot* of reading once you start graduate school, so you might as well get used to it.) This is where advisors first enter the scene: faculty members ought to be willing to talk to undergraduates and help them find out more about research areas and graduate schools. Try to get involved in research: ask professors and TAs whether they need someone to work on an ongoing project, or start an independent research project, with guidance from a faculty member.

Contact faculty members and graduate students at the schools you're interested in. Tell them about your background and interests and ask them what research projects they're working on. A good way to do this is via electronic mail if possible – e-mail is much easier and quicker to respond to than a paper letter. A good advisor will be willing to answer these kinds of inquiries (al-

though if they're busy they may give you only a brief answer or point you towards a graduate student – you'll have to use your intuition to decide whether they're brushing you off or just busy). If you can't get any answer at all, consider that that individual might not end up being a very accessible advisor. Asking these questions will help you narrow down your choices and may increase your chances of admission if the professors you contact become interested in working with you.

Your best bet is to find a school where there are at least two faculty members you'd be interested in working with. That way, if one doesn't work out, or is too busy to take on a new student, you have a fallback position. It's also important to most people to feel comfortable with the community of graduate students. It pays to talk to some of the graduate students (both junior and senior) to find out how they like it, which advisors are good, and what kinds of support (financial and psychological) are available. Because there are so many students applying to each school, even highly qualified applicants are often rejected. You should apply to a range of programs – and don't take it personally if you do get rejected by some of them.

You can increase your chances of getting into graduate school by getting good grades, especially in upper division classes in your area of interest, having a broad background in your field and in related fields (for example, psychology classes are useful for AI students), getting a high score on the GRE if required, developing good relationships with your professors and work managers (this is very important for getting good recommendations), working on a research project, and having a clear sense of what you want to work on (although it's always all right to change your mind later). Also, it's a good idea to start thinking early about sources of funding: apply for an NSF fellowship and ask your advisor or department office about other fellowships.

3 Doing Research

For many new graduate students, graduate school is unlike anything else they've done.

Sometimes it's hard to know exactly what it is you're supposed to be learning. Yes, you have to complete a dissertation, but how do you start? What should you spend your time doing?

Graduate school is a very unstructured environment in most cases. Graduate students typically take nine hours or less of coursework per semester, especially after the second year. For many, the third year – after coursework is largely finished and preliminary exams have been completed – is a very difficult and stressful period. This is when you're supposed to find a thesis topic, if you're not one of the lucky few who has already found one. Once you do find a topic, you can expect two or more years until completion, with very few landmarks or milestones in sight.

The following sections talk about the day-to-day process of doing research, criticism and feedback, working on the thesis, and financial support for research.

3.1 The Daily Grind

Being a good researcher involves more than “merely” coming up with brilliant ideas and implementing them. Most researchers spend the majority of their time reading papers, discussing ideas with colleagues, writing and revising papers, staring blankly into space – and, of course, having brilliant ideas and implementing them.

Section 5 discusses the process and importance of becoming part of a larger research community, which is a critical aspect of being a successful researcher. This section contains ideas on keeping track of where you're going, and where you've been, with your research, staying motivated, and how to spend your time wisely.

Keeping a journal of your research activities and ideas is very useful. Write down speculations, interesting problems, possible solutions, random ideas, references to look up, notes on papers you've read, outlines of papers to write, and interesting quotes. Read back through it periodically. You'll notice that the bits of random thoughts start to come together and form a pattern, often turning into a research project or even a thesis topic. I was surprised, looking

back through my journal as I was finishing up my thesis, how early and often similar ideas had cropped up in my thinking, and how they gradually evolved into a dissertation.

You'll have to read a lot of technical papers to become familiar with any field, and to stay current once you've caught up. You may find yourself spending over half of your time reading, especially at the beginning. This is normal. It's also normal to be overwhelmed by the amount of reading you think you “should” do. Try to remember that it's impossible to read everything that might be relevant: instead, read selectively. When you first start reading up on a new field, ask your advisor or a fellow student what the most useful journals and conference proceedings are in your field, and ask for a list of seminal or “classic” papers that you should definitely read. For AI researchers, a useful (if slightly outdated) starting point is Agre's [1982] summary of basic AI references. Similar documents may exist for other research areas – ask around. Start with these papers and the last few years of journals and proceedings.

Before bothering to read *any* paper, make sure it's worth it. Scan the title, then the abstract, then – if you haven't completely lost interest already – glance at the introduction and conclusions. (Of course, if your advisor tells you that this is an important paper, skip this preliminary step and jump right in!) Before you try to get all of the nitty-gritty details of the paper, skim the whole thing, and try to get a feel for the most important points. If it still seems worthwhile and relevant, go back and read the whole thing. Many people find it useful to take notes while they read. Even if you don't go back later and reread them, it helps to focus your attention and forces you to summarize as you read. And if you do need to refresh your memory later, rereading your notes is much easier and faster than reading the whole paper.

A few other points to keep in mind as you read and evaluate papers:

- Make sure the ideas described really worked (as opposed to just being theoretically valid, or tested on a few toy examples).

- Try to get past buzzwords: they may sound good, but not mean much. Is there substance and an interesting idea underneath the jargon?
- To really understand a paper, you have to understand the motivations for the problem posed, the choices made in finding a solution, the assumptions behind the solution, whether the assumptions are realistic and whether they can be removed without invalidating the approach, future directions for research, what was actually accomplished or implemented, the validity (or lack thereof) of the theoretical justifications or empirical demonstrations, and the potential for extending and scaling the algorithm up.

Keep the papers you read filed away so you can find them again later, and set up an online bibliography (BibTeX is a popular format, but anything consistent will do). I find it useful to add extra fields for keywords, the location of the paper (if you borrowed the reference from the library or a friend), and a short summary of particularly interesting papers. This bibliography will be useful for later reference, for writing your dissertation, and for sharing with other graduate students (and eventually, perhaps, advisees).

3.2 Staying Motivated

At times, particularly in the “middle years,” it can be very hard to maintain a positive attitude and stay motivated. Many graduate students suffer from insecurity, anxiety, and even boredom. First of all, realize that these are normal feelings. Try to find a sympathetic ear – another graduate student, your advisor, or a friend outside of school. Next, try to identify why you’re having trouble and identify concrete steps that you can take to improve the situation. To stay focused and motivated, it often helps to have organized activities to force you to manage your time and to do something every day. Setting up regular meetings with your advisor, attending seminars, or even extracurricular activities such as sports or music can help you to maintain a regular schedule.

Chapman [1988] enumerates a number of “immobilizing shoulds” that can make you feel so guilty and unworthy that you stop making progress. Telling yourself that you *should* have a great topic, that you *should* finish in n years, that you *should* work 4, or 8, or 12 hours a day isn’t helpful for most people. Be realistic about what you can accomplish, and try to concentrate on giving yourself positive feedback for tasks you do complete, instead of negative feedback for those you don’t.

Setting daily, weekly, and monthly goals is a good idea, and works even better if you use a “buddy system” where you and another student meet at regular intervals to review your progress. Try to find people to work with: doing research is much easier if you have someone to bounce ideas off of and to give you feedback.

Breaking down any project into smaller pieces is always a good tactic when things seem unmanageable. At the highest level, doing a master’s project before diving into a Ph.D. dissertation is generally a good idea (and is mandatory at some schools). A master’s gives you a chance to learn more about an area, do a smaller research project, and establish working relationships with your advisor and fellow students.

The divide-and-conquer strategy works on a day-to-day level as well. Instead of writing an entire thesis, focus on the goal of writing a chapter, section, or outline. Instead of implementing a large system, break off pieces and implement one module at a time. Identify tasks that you can do in an hour or less; then you can come up with a realistic daily schedule. If you have doubts, don’t let them stop you from accomplishing something – take it one day at a time. Remember, every task you complete gets you closer to finishing.

3.3 Getting to the Thesis

The hardest part of getting a Ph.D. is, of course, writing the dissertation. The process of finding a thesis topic, doing the research, and writing the thesis is different from anything most students have done before. If you have a good advisor

and support network, you'll be able to get advice and help in setting directions and goals. If not, you may need to be more independent. If this is the case, don't just isolate yourself from the world: try to go out and find the resources and support you need from professors, other graduate students, mailing lists, friends, family, and publications like this one.

3.3.1 Finding an Advisor

Finding the right advisor can help you immeasurably in successfully completing a thesis. You should ideally have selected the schools you applied to by identifying faculty members you'd like to work with. If not, start looking around as early as possible. Of course, the ideal advisor will be in the area you're interested in working in, and will actively be doing high-quality research and be involved in and respected by the research community.

Read research summaries by faculty members (which are usually published by the department), go to talks they give, and attend or audit courses given by professors you might be interested in working with. Talk to other graduate students and recent graduates. Ask them how their relationships with their advisors are/were, how quickly the advisor's students graduate, and how successful (well recognized, high-quality) their research is. What kinds of relationships do they have – frequent interactions, collaborative work, encouraging independence? handing out topics or helping students to create individual research areas, or a more hands-off style?

Other things to find out about potential advisors:

- What is the average time their Ph.D. students take to finish their degrees? What is the dropout rate for their students?
- How long have they been on the faculty? There are advantages and disadvantages to being one of the first members of a new research group. On the positive side, you often have more freedom to choose your research topic and to influence the direction of

the group's research. On the negative side, you may be more isolated (since there won't be older graduate students in the group), and your advisor won't have as much experience.

A good advisor will serve as a mentor as well as a source of technical assistance. A mentor should provide, or help you to find, the resources you need (financial, equipment, and psychological support); introduce you and promote your work to important people in your field; encourage your own interests, rather than promoting their own; be available to give you advice on the direction of your thesis and your career; and help you to find a job when you finish. They should help you to set and achieve long-term and short-term goals.

Once you identify one or more potential advisors, get to know them. Introduce yourself and describe the area you're interested in. Attend their research group meetings if they hold them regularly. Give them a copy of a research proposal if you have a good idea of what you want to work on, and ask for comments. Ask whether they have any TA or RA positions available, or if there are any ongoing research projects that you could get involved with. Read their published papers, and the work of their students. Drop by during office hours and ask questions or make comments. Offer to read drafts of papers – and do more than just proofread (see Section 3.4).

The type of relationship that each student needs with an advisor will be different. Some students prefer to be given more direction, to have frequent contact, and to be “checked up on.” Others are more independent. Some may need contact but be self-conscious about asking for it. Other things that vary include what kinds of feedback is preferred (lots of “random” ideas vs. very directed feedback (pointers)), working individually vs. in groups, working on an established research project vs. a new, independent effort; working in the same area as your advisor or doing an “outside” thesis.

You may find that your thesis advisor doesn't always give you all of the mentoring that you need. Multiple mentors are common and useful;

they may include other faculty members in your department or elsewhere, senior graduate students, or other colleagues (see Section 5.3). You may want to seriously consider changing thesis advisors if your advisor is inaccessible or disinterested, gives you only negative feedback, doesn't have the technical background to advise you on your thesis, or harasses you (see Section 7).

The most important thing is to ask for (i.e., demand politely) what you need.

3.3.2 Finding a Thesis Topic

Doing a master's project is often a good idea (and is required by some schools). Although choosing an appropriately scaled-down topic may be difficult, having the ideal topic is also less important, since you will have the chance to move on after only a year or so. If you have a good idea of what you want to do your Ph.D. dissertation on, choosing a master's project that will lead into the dissertation is wise: you will get a head start on the Ph.D., or may decide that you're not interested in pursuing the topic after all (saving yourself a lot of work and grief farther down the road).

A good source of ideas for master's projects (and sometimes for dissertation topics) is the future work section of papers you're interested in. Try developing and implementing an extension to an existing system or technique.

Generally speaking, a good Ph.D. thesis topic is interesting to you, to your advisor, and to the research community. As with many aspects of graduate school, the balance you find will depend at least in part on the relationship you have with your advisor. Some professors have well defined long-term research programs and expect their students to contribute directly to this program. Others have much looser, but still related ongoing projects. Still others will take on anyone with an interesting idea, and may have a broad range of interesting ideas to offer their students. Be wary of the advisor who seems willing to let you pursue any research direction at all. You probably won't get the technical support you need, and they may lose interest in you

when the next graduate student with a neat idea comes along.

If you pick a topic that you're not truly interested in simply because it's your advisor's pet area, it will be difficult to stay focused and motivated – and you may be left hanging if your advisor moves on to a different research area before you finish. The same is true for choosing a topic because of its marketability: if you're not personally excited about the topic, you'll have a harder time finishing and a harder time convincing other people that your research is interesting. Besides, markets change more quickly than most people finish dissertations.

In order to do original research, you must be aware of ongoing research in your field. Most students spend up to a year reading and studying current research to identify important open problems. However, you'll never be able to read everything that might be relevant – and new work is always being published.

Try to become aware and stay aware of directly related research – but if you see new work that seems to be doing exactly what you're working on, don't panic. It's common for graduate students to see a related piece of work and think that their topic is ruined. If this happens to you, reread the paper several times to get a good understanding of what they've really been accomplished. Show the paper to your advisor or someone else who's familiar with your topic and whose opinions you respect. Introduce yourself to the author at a conference or by e-mail, and tell them about your work. By starting a dialogue, you will usually find that their work isn't quite the same, and that there are still directions open to you. You may even end up collaborating with them. Good researchers welcome the opportunity to interact and collaborate with someone who's interested in the same problems they are.

To finish quickly, it's usually best to pick a narrow, well defined topic. The downside of this approach is that it may not be as exciting to you or to the research community. If you're more of a risk-taker, choose a topic that branches out in a new direction. The danger here is that it can be

difficult to carefully define the problem, and to evaluate the solution you develop. If you have a topic like this, it helps a lot to have an advisor or mentor who is good at helping you to focus and who can help you maintain a reasonably rigorous approach to the problem.

In the extreme case, if your topic is so out of the ordinary that it's unrelated to anything else, you may have difficulty convincing people it's worthwhile. Truly innovative research is, of course, exciting and often pays back in recognition from the research community – or you could just be out in left field. If you have a far-out topic, be sure that people are actually *interested* in it, or you'll never be able to “sell” it later, and will probably have trouble getting your work published and finding a job. In addition, it will be hard to find colleagues who are interested in the same problems and who can give you advice and feedback.

In any case, a good topic will address important issues. You should be trying to solve a real problem, not a toy problem (or worse yet, no problem at all); you should have solid theoretical work, good empirical results or, preferably, both; and the topic will be connected to – but not be a simple variation on or extension of – existing research. It will also be significant yet manageable. Finding the right size problem can be difficult. One good way of identifying the right size is to read other dissertations. It's also useful to have what Chapman [1988] calls a “telescoping organization” – a central problem that's solvable and acceptable, with extensions and additions that are “successively riskier and that will make the thesis more exciting.” If the gee-whiz additions don't pan out, you'll still have a solid result.

Remember that a thesis is only a few years of your work, and that – if all goes well – your research career will continue for another 30 or 40. Don't be afraid to leave part of the problem for future work, and don't compare yourself to senior researchers who have years of work and publications to show for it. (On the other hand, if you identify too much future work, your thesis won't look very exciting by comparison.) Graduate students often pick overly ambitious topics

(in theory, your advisor will help you to identify a realistic size problem). Don't overestimate what other people have done. Learn to read between the lines of grandiose claims (something else a good advisor will help you to do).

Some schools may require that you write a thesis proposal. Even if they don't, this is a good first step to take. It forces you to define the problem, outline possible solutions, and identify evaluation criteria; and it will help you to get useful feedback from your advisor and other colleagues. Writing a good thesis proposal will take up to several months, depending on how much background work and thinking you've already done in the process of choosing the topic.

The proposal should provide a foundation for the dissertation. First, you must circumscribe the problem and argue convincingly that it needs to be solved, and that you have a methodology for solving it. You must identify and discuss related work: has this problem been addressed before? What are the shortcomings of existing work in the area, and how will your approach differ from and be an improvement over these methods?

Present your ideas for solving the problem in as much detail as possible, and give a detailed plan of the remaining research to be done. The proposal should include, or be structured as, a rough outline of the thesis itself. In fact, unless your final topic differs significantly from your proposed topic (which many do), you may be able to reuse parts of the proposal in the thesis.

You will probably have to take an oral exam in which you present and/or answer questions about your proposal. Be sure that your committee members are as familiar as possible with your work beforehand. Give them copies of the proposal, and talk to them about it. During the exam, don't panic if you don't know the answer to a question. Simply say, “I'm not sure” and then do your best to analyze the question and present possible answers. Your examining committee wants to see your analytical skills, not just hear canned answers to questions you were expecting. Give a practice talk to other students and faculty members. Remember: you know

more about your thesis topic than your committee; you're teaching *them* something for a change.

3.3.3 Writing the Thesis

Graduate students often think that the thesis happens in two distinct phases: doing the research, and writing the dissertation. This may be the case for some students, but more often, these phases overlap and interact with one another. Sometimes it's difficult to formalize an idea well enough to test and prove it until you've written it up; the results of your tests often require you to make changes that mean that you have to go back and rewrite parts of the thesis; and the process of developing and testing your ideas is almost never complete (there's always more that you *could* do) so that many graduate students end up "doing research" right up until the day or two before the thesis is turned in.

The divide-and-conquer approach works as well for writing as it does for research. A problem that many graduate students face is that their only goal seems to be "finish the thesis." It is essential that you break this down into manageable stages, both in terms of doing the research and when writing the thesis. Tasks that you can finish in a week, a day, or even as little as half an hour are much more realistic goals. Try to come up with a range of tasks, both in terms of duration and difficulty. That way, on days when you feel energetic and enthusiastic, you can sink your teeth into a solid problem, but on days when you're run-down and unmotivated, you can at least accomplish a few small tasks and get them off your queue.

It also helps to start writing at a coarse granularity and successively refine your thesis. Don't sit down and try to start writing the entire thesis from beginning to end. First jot down notes on what you want to cover; then organize these into an outline (which will probably change as you progress in your research and writing). Start drafting sections, beginning with those you're most confident about. Don't feel obligated to write it perfectly the first time: if you can't get a paragraph or phrase right, just write *something*

(a rough cut, a note to yourself, a list of bulleted points) and move on. You can always come back to the hard parts later; the important thing is to make steady progress.

When writing a thesis, or any technical paper, realize that your audience is almost guaranteed to be less familiar with your subject than you are. Explain your motivations, goals, and methodology clearly. Be repetitive without being boring, by presenting your ideas at several levels of abstraction, and by using examples to convey the ideas in a different way.

Having a "writing buddy" is a good idea. If they're working on their thesis at the same time, so much the better, but the most important thing is that they be willing to give you feedback on rough drafts, meet regularly to chart your progress and give you psychological support, and preferably that they be familiar enough with your field to understand and review your writing.

3.4 Getting Feedback

To be successful at research, it is essential that you learn to cope with criticism, and even that you actively seek it out. Learn to listen to valid, constructive criticism and to ignore destructive, pointless criticism (after finding any pearls of wisdom that may be buried in it).

In order to get feedback, you have to present your ideas. Write up what you're working on, even if you're not ready to write a full conference or journal paper, and show it to people. Even for pre-publishable papers, write carefully and clearly, to maximize your chances of getting useful comments (and of having people read what you wrote at all).

Give presentations at seminar series at your university, at conferences, and at other universities and research labs when you get the chance. Your advisor should help you find appropriate forums to present your work and ideas. Many fields have informal workshops that are ideal for presenting work in progress.

Attend conferences and talk about your research. When you meet someone new and they

ask you what you're working on, seize the opportunity. Don't just say "I'm doing my thesis on foobar applications of whatsis algorithms" – tell them as much as they're willing to listen to. You should have 30-second, 2-minute, 5-minute and 10-minute summaries of your thesis ready at a moment's notice (but not memorized word-for-word; nobody wants to listen to a canned speech).

Talking to other people will help you to realize which aspects of your research are truly different and innovative, how your work fits into the current state of your field and where it's going, and which aspects of your work are harder to sell (and, therefore, which aspects you need to think more about justifying).

Giving feedback to other students and colleagues is useful for many reasons. First, it helps you to polish your critical skills, which are helpful both in understanding other people's work and in evaluating your own. Second, it helps you to build a network of people who will be your colleagues for years to come. Finally, if you give useful feedback, those people will be more likely to make an effort to do the same for you.

It will be helpful (to you and to the person whose paper you're reviewing) to organize comments on a paper in descending order of abstraction: high-level content-oriented comments, mid-level stylistic and presentation comments, and low-level nitpicky comments on syntax and grammar. Try to keep your comments constructive ("this would read better if you defined X before introducing Y") rather than destructive ("this is nonsense").

You'll want to read a paper at least twice – once to get the basic ideas, then a second time to mark down comments. High-level comments describing your overall impression of the paper, making suggestions for organization, presentation and alternative approaches to try, potential extensions, and relevant references are generally the most useful and the hardest to give. Low-level comments are more appropriate for a paper that is being submitted for publication than for an unpublished paper such as a proposal or description of preliminary research.

3.5 Getting Financial Support

Most graduate students (at least in the natural sciences) have a source of financial support that pays their tuition and a small living stipend. Although nobody ever got rich being a graduate student, you probably won't starve either. Sources of funding include fellowships (from NSF, universities, foundations, government agencies, and industry), employer support, research assistantships (i.e., money from a faculty member's research grant) and teaching assistantships.

Start looking for money early. Many schools arrange support in the form of an RA or TA position in the first year, but after that, you're on your own. Deadlines for applications vary, and if you miss one, you'll probably have to wait another year. After you apply, it can take six months or so to review the applications and several more months to actually start receiving money.

Ask faculty members (especially your advisor, who should be helping you to find support or providing support out of his or her grant money), department administrators, and fellow graduate students about available funding. Go to your university's fellowship office or its equivalent, and look through the listings in *The Annual Register of Grant Support*, *The Grant Register*, *The Chronicle of Higher Education*, and *Foundation Grants to Individuals*. Look into NSF grants (there are several different programs). Take advantage of your status as a woman or minority if you are one (this may be the only time when it actually is an advantage). Most universities have fellowship programs that may be administered through individual departments or may be campus-wide.

If you haven't yet begun actively doing research, getting an RA position from a faculty member may be a good way to become involved in a research project. Working on an existing research project by maintaining or developing hardware or software, writing reports, and running experiments will give you a feel for what it's like to do research – and you may even find a thesis topic. Ask around to see what's available,

and go talk to professors whose work you find interesting.

For a research grant or fellowship, you will probably have to write a proposal, so the more you've thought about potential thesis topics, the better off you'll be. You may need to tailor your proposal to the interests and needs of the particular funding agency or program you're applying to, but stick to something you know about and are sincerely interested in.

Write for a general audience, since the people reviewing your application may not be in the same field. Emphasize your goals and why the project you propose to work on is important. Talk as much as you can about how you're going to solve the problem, and be sure that your proposed solution will satisfy the goals you've set forth. Follow the rules for format, page layout and length, or your application may not even be reviewed.

4 Advice for Advisors

In order to be a good advisor, you have to relate to your graduate students as individuals, not just as anonymous research assistants or tickets to tenure and co-authored publications. Work with all of your graduate students, not just those whom you feel most comfortable with, or who are interested in the problems you're most excited about. Try to get to know your students personally and professionally. Help them to identify their strengths and weaknesses, to build on the former, and to work on overcoming the latter. Give them honest evaluations of their work and performance: don't just assume that they know how they're doing and what you think of them.

Read this paper and others like it with an eye towards discovering which aspects of the graduate experience your students may be having trouble with, or may not realize the importance of. Try to see the experience from their perspective, which will be different for each student, because each student has a different background and different talents and goals.

The roles of an advisor include:

- Guiding students' research: helping them to select a topic, write a research proposal, perform the research, evaluate it critically, and write the dissertation.
- Getting them involved in the wider research community: introducing them to colleagues, collaborating on research projects with them, funding conference travel, encouraging them to publish papers, nominating them for awards and prizes.
- Finding financial support: providing research assistantships or helping them to find fellowships, and finding summer positions.
- Finding a position after graduation: helping them to find and apply for postdoctoral positions, faculty positions, and/or jobs in industry; supporting their applications with strong recommendations; and helping them to make contacts.

Although guiding your students' research is normally viewed as the central task of an advisor, the other roles are also critical to their long-term success. Section 5.3 contains advice for students on networking. You can help them in this process by funding and encouraging travel to conferences and paper publication, and by introducing them and talking about their research to colleagues.

4.1 Interacting With Students

Especially for a new advisor, setting the right tone for student interactions is a difficult task. Different students respond best to different approaches – and, of course, different advisors have different personal styles. Some of the tradeoffs that have to be made in each advisor-student relationship are:

- Amount of direction: self-directed/hands-off vs. “spoon-feeding” topics and research projects.
- Personal interactions and psychological support: do they want advice on career, family, and the like? Are you willing and able to give it, or to find someone else to advise them?

- Amount and type of criticism: general directions vs. specific suggestions for improvement.
- Frequency of interaction: daily vs. once a semester.

It helps to establish regular meeting times and to discuss expectations (both yours and your students') about what can and should be accomplished during these meetings. Encourage them to develop relationships with other faculty members, students, and colleagues, to get a different perspective and to get feedback you may not be able to give.

To improve the atmosphere of your interactions:

- Meet over lunch or coffee to make interactions more relaxed and less stressful.
- Strive to maintain an open, honest relationship. Respect your students as colleagues.
- Tell them if you think they're asking for too much or too little time or guidance.

Advisors should be aware of both long-term and short-term needs. What should the student's goals over the next few years be? Help your student identify ways that the two of you – as a team – can meet these goals. Advise the student on the criteria for a successful qualifying exam, thesis proposal, and dissertation. Help prepare the student for a future research career.

In the short term, a good advisor will work with students to set priorities and to find a balance between doing research, reading, writing, satisfying TA and RA duties, publishing, and coursework. Although advisors may not be able to give advice on all administrative aspects of graduate school, they should at least know the appropriate people to refer students to for assistance with degree requirements, funding, and so on.

When you meet with your students, pay attention to them. Try to help them to identify their interests, concerns, and goals, not just how can they meet what *you* see as good interests,

concerns, and goals. Know what they're working on, and what you discussed last time. Take notes during meetings and review them if you have to.

Give them productive feedback, not just a noncommittal "ok, sure" or a destructive "why on earth do you want to do that?" Remember that your students are still learning. If you tell them that a problem they're interested in has already been explored by Professor X, make sure you follow up with a reference that they have access to, and a discussion as to whether the problem remains a worthwhile area to work on, or whether there are new open issues raised by Professor X's work, at the next meeting.

When reviewing a student's paper or proposal, write comments on the paper itself: verbal comments aren't as useful. Give the feedback promptly, or it won't be much help. See the section on feedback (Section 3.4) for suggestions about giving useful comments. Don't just wait until they hand you something to read: insist on written drafts of proposals, papers, etc. Help them develop their rough ideas into publishable papers. Give them specific, concrete suggestions for what to do next, especially if they seem to be floundering or making little progress.

Advisor-student relationships can break down if the advisor is setting goals that are too high or too low, or if the advisor is exploiting the student to meet the advisor's needs, not the student's. In my opinion, it is never appropriate to develop an intimate relationship with one of your own students. If this should happen, you should not continue to advise them (whether the relationship continues or not).

Encourage your students to choose a topic that you're *both* interested in and that you're knowledgeable about (or very interested in learning more about). Make sure that they have the appropriate background to understand the problem, and that the methodology and solution they identify are appropriate and realistic. Give them pointers to useful references and help them find them (this can be a mysterious, difficult process for graduate students). Make sure they're aware of other researchers and labs who are doing sim-

ilar work, and if possible, arrange for them to visit these labs or meet the researchers at seminars or conferences.

Women faculty often feel obligated to mentor every woman student in the department, attend every committee meeting, and get involved in every debate, whether they want to or not. While you can't solve all of the problems in the world, you can at least make a difference by giving other women (and men, for that matter) the sense that you do care, and that you think women's issues are important, even if you don't have time (or the inclination) to get involved with every problem.

5 Becoming Part of the Research Community

One of the most important things a graduate student should do is to become established as part of the research community. Your advisor can help with this process by funding conference travel, encouraging you to publish research results early, collaborating on joint publications, introducing you to colleagues, and promoting your work.

In turn, you can make yourself more visible by participating in conferences and workshops, publishing papers on your work, and meeting and maintaining contact with colleagues.

5.1 Attending Conferences

Attending conferences and workshops is valuable whether you present a paper or not. Some of the reasons to do so are:

- You'll meet people and have a chance to discuss your ideas and to hear theirs.
- You'll get a good sense of what the current state of research is, and will learn more about how to write conference papers and give talks (sometimes by counterexample).
- You'll probably realize that your ideas are more significant, relatively speaking, than

you thought. A common reaction is "I could write a better paper than this!"

If you're giving a talk you'll gain even more visibility, and will have an opportunity to make an impression on other researchers. Some tips for preparing your talk to make this impression as positive as possible:

- Give a practice talk, especially if you tend to get stage fright. Be sure to invite people who will give you constructive, but useful, feedback.
- Make sure your talk fits in the time slot allocated. There's nothing worse than a speaker who rushes through the last ten slides, or skips from the middle of the talk to the conclusion. A good rough rule is to allocate 2-3 minutes per slide, on average.
- It's better to be somewhat abstract than to get bogged down in technical details – but be sure you give enough detail to make a convincing case. Your paper should fill in the missing details, so that people can read it to get a more in-depth understanding. Know your audience: you'll have to give more background to a general audience, and more technical detail to audiences that are very familiar with the field of research you're discussing.
- Use examples and pictures to illustrate and clarify your ideas.
- Learn by observation: try to imitate qualities of talks that you like, and avoid things that other speakers do that bother you.
- Talk about your ideas informally whenever you get the chance, so that the talk will come more naturally and, hopefully, you'll have a chance to respond to and think about questions that might get asked at the talk.
- Make sure your slides are readable and as simple as possible. Never put up a slide with tiny text and say "I know you can't read this, but..."

- Try to relax. Don't read from a script or word-for-word from your slides, and don't talk too fast. Be confident: you know more about your work (flaws and all) than anyone else.

5.2 Publishing Papers

Publishing your ideas is important for several reasons: it gives you a source of feedback from people who read your papers; it establishes you as a member of the research community (useful for getting a job down the line); and it forces you to clarify your ideas and to fit them in the context of the current state of research in your field.

There are two key properties of a good paper: significant content – original, important ideas that are well developed and tested – and good writing style. The degree to which the paper's content has to be “significant” depends on where you're submitting it. Preliminary ideas and work in progress are more suitable for a workshop or symposium; well developed, extensively tested ideas are more appropriate for a journal. One way to decide where your paper should be submitted is to read papers in potentially appropriate publications (last year's conference proceedings; current journal issues). Another method to show a draft or outline of the paper to your advisor or other colleagues and ask their advice.

if you have a great idea, but present it poorly, your paper probably won't be accepted. Be sure you know what the point of the paper is, and state it clearly and repeatedly. The same goes for the key technical ideas. Don't make the reader work to figure out what's important – tell them explicitly. Otherwise, they might get it wrong, if they bother to finish reading the paper at all. State the problem you're addressing, why it's important, how you're solving it, what results you have, how other researchers have addressed the same or similar problems, and why your method is different or better.

Write for the audience that you expect to read the paper, just as you would plan a talk. Give more background for general audiences, less

background and more technical detail for specialized audiences. Use a running example if possible, especially if your paper is dense with equations and algorithms.

Don't try to put every idea in your thesis into one conference paper. Break it down into pieces, or write one or two longer journal articles.

As you refine your ideas, you can re-publish in new forms, but be sure you're adding new material, not just rehashing the same ideas. Some papers start as short workshop papers, evolve into conference papers, and eventually – with the addition of detailed empirical results or formal proofs – become journal articles. It's usually okay to publish the same or substantially similar papers in multiple workshops, but papers for conferences and journals generally have to be original, unpublished work.

It is critical that any paper you plan to submit be read by someone else first, if only to check for typos, grammatical errors, and style. A good reviewer will give you feedback on the organization and content of the paper as well (see Section 3.4). The more tightly refereed the publication you're submitting to, the more trouble you should go to to have it pre-reviewed. For a workshop paper, having your advisor read it over is probably enough. For a refereed conference, have one or two other graduate students read it as well. For a journal paper, you should probably find researchers who are active in the field, preferably at other institutions (to give breadth), read it over and give you comments. This is where the network of colleagues you should build (Section 5.3) comes in handy.

If your paper is rejected, keep trying! Take the reviews to heart and try to rewrite the paper, addressing the reviewer's comments. You'll get more substantial and useful reviews from journals than conferences or workshops. Often a journal paper will be returned for revisions; usually a conference paper will just be accepted or rejected outright. After reading the review the first time, put it aside. Come back to it later, reading the paper closely to decide whether the criticisms were valid and how you can address them. You will often find that reviewers make

criticisms that are off-target because they misinterpreted some aspect of your paper. If so, don't let it get to you – just rewrite that part of your paper more clearly so that the same misunderstanding won't happen again. It's frustrating to have a paper rejected because of a misunderstanding, but at least it's something you can fix. On the other hand, criticisms of the content of the paper may require more substantial revisions – rethinking your ideas, running more tests, or redoing an analysis.

5.3 Networking

One of the most important skills you should be learning in graduate school is how to “network.” Breaking into the research community requires attending conferences, meeting established researchers, and making yourself known. Networking *is* a learned skill, so you shouldn't expect to be an expert at it immediately; but it is also a skill that you can, and should, learn in order to be a successful member of the research community.

Just going to conferences and standing in the corner isn't enough. Especially if you're not normally an outgoing person, you have to make a conscious effort to meet and build relationships with other researchers. Presenting papers is a good way to do this, since people will often approach you to discuss your presentation. Introducing yourself to people whose presentations you found interesting, and asking a relevant question or describing related research you're doing, is also a good way to meet people.

You should talk about your research interests every chance you get. (But be sure to spend some time listening, too: you'll learn more this way, and people will feel that your conversations are a two-way street.) Have summaries of your work of various lengths and levels of detail mentally prepared, so that you can answer the inevitable “So what are you working on?” intelligently and clearly. If someone expresses an interest in your work, follow up! Send them e-mail talking about new ideas or asking questions; send them drafts of papers; ask them for drafts of their

papers and send them comments. (If you do this, they'll be sure to remember you!) Bring business cards with your e-mail address to conferences to help new acquaintances jog their memory.

Maintain the relationships you form via e-mail, and by re-establishing contact at each workshop or conference you attend. If you work at it, and use your initial acquaintances to meet new people, you'll find that your “network” grows rapidly.

Sometimes these contacts will grow into opportunities to do collaborative research. Seize these opportunities: you will meet more people, often become exposed to new methods of doing research or new subfields within your research area, and the responsibility you feel towards your collaborator may give you more of an incentive to stay motivated and keep accomplishing something.

Other professional activities can bring you into the research network as well: volunteer for program committees, send your resume to a book review editor, offer to give seminars at other universities, write conference and workshop papers and send them to people you've met or would like to meet, or organize a workshop on your subfield at a larger conference. Mentoring junior graduate students and undergraduates is a good investment in the long run (besides providing them a valuable service and making you feel useful and knowledgeable).

Finding specific mentors can be very useful. Especially if you feel that you are isolated at your institution, having a colleague at another institution who can give you advice, feedback on drafts of papers, and suggestions for research directions can be extremely valuable.

6 All Work and No Play...

Finding a balance between work, play, and other activities isn't easy. Different people will give you very different advice. Some people say you should be spending eighty or ninety percent of your waking hours working on your thesis. Others (myself included) think that this is unrealistic and unhealthy, and that it's important for you

mental and physical health to have other active interests.

If you have a family, you will have to balance your priorities even more carefully. Graduate school isn't worth risking your personal relationships over; be sure that you save time and energy to focus on the people who matter to you.

One of the keys to balancing your life is to develop a schedule that's more or less consistent. You may decide that you will only work during the days, and that evenings are for your hobbies. Or you might decide that afternoons are for socializing and exercising, and work late at night. I decided very early on in graduate school that weekends were for me, not for my thesis, and I think it helped me to stay sane.

Many graduate students hit the doldrums around the end of the second or beginning of the third year, when they're finishing up their coursework and trying to focus in on a thesis topic. Sometimes this process can take quite a while. Try to find useful, enjoyable activities that can take your mind off of the thesis. Sing in a choir, learn a foreign language, study the history of ancient Greece, garden, or knit. If you schedule regular activities (rehearsals, tennis lessons), you will probably find it easier to avoid drifting aimlessly from day to day.

In the final push to finish your thesis, though, you will almost certainly have less time for social activities than you used to. Your friends may start to make you feel guilty, whether they intend to or not. Warn them in advance that you expect to turn down lots of invitations, and it's nothing personal – but you need to focus on your thesis for a while. Then you'll be all done and free as a bird! (Until the next phase of your life starts...)

7 Issues for Women

Although this paper started out from a discussion about the problems women face in graduate school, it has evolved into something that I think is relevant for everyone, not just women. This is not to say, however, that there aren't special problems faced by women.

In many cases, women and men face the same

obstacles in graduate school, but react differently to them. For women, the additional factors that are sometimes (but not always) present include isolation, low self-esteem, harassment and discrimination, unusual time pressures arising from family responsibilities, lack of a support network, and lack of relevant experience. Having an un-supportive advisor can thus become much more of a problem for women than for men. I hope that to some extent, this paper will help both women and advisors of women to provide the supportive, positive environment that all graduate students deserve.

Part of the reason that I changed the focus of the paper is that there have been many articles written recently on the subject of women scientists and women graduate students. These include [Spertus, 1992, Toth, 1988, Hall and Sandler, b, Hall and Sandler, c, Hall and Sandler, a, Sandler and Hall, 1986, NSF, 1989, Leveson, 1989, Strok, 1992]; [McKay, 1988] talks about issues relevant for minority faculty members, many of which pertain to minority graduate students. The *systems* mailing list is an electronic resource for women in computer science; send e-mail to `systems-request@pa.dec.com` for more information.

8 Conclusions

In addition to the papers I have cited directly in the article, I found a variety of other resources to be useful, and have included them in the References section.

The UC Berkeley *Graduate* is a newsletter published by the UC Berkeley Graduate Division with articles of general interest to graduate students. I found this publication very informative both during graduate school and while writing this article. A number of particularly interesting articles are included in the References section.

Several articles ([Bundy *et al.*, 1986, Bental, 1992, Chapman, 1988]) give general advice on graduate school and doing research. Guidelines and suggestions for reviewing papers are given in [Smith, 1990] and [Shriver, 1990].

A number of articles on writing proposals

and successfully applying for research grants are available ([Somerville, 1982, White, 1983, White, 1975, Lefferts, 1978]).

Graduate school is not an easy process, and too many students are thwarted and intimidated by unsupportive or unskilled advisors, lack of knowledge about what graduate school is all about, inflexible bureaucracies, and a myriad of other obstacles. I have tried to give advice that graduate students and caring advisors can use to lessen some of these obstacles.

References

- [Agre, 1982] Philip E. Agre. What to read: A biased guide to AI literacy for the beginner. Technical Report Working Paper 239, MIT AI Lab, November 1982.
- [Bental, 1992] Diana Bental. Thesis prevention: Advice to phd supervisors. *AISB Quarterly No. 80 (Newsletter of the Society for the Study of Artificial Intelligence and Simulation of Behaviour)*, pages 58–60, Summer 1992. (Published under the alias ‘The Siblings of Perpetual Prototyping’).
- [Bundy *et al.*, 1986] Alan Bundy, Ben du Boulay, Jim Howe, and Gordon Plotkin. The researchers’ bible. Technical Report DAI Teaching Paper No. 4, Dept. of Artificial Intelligence, University of Edinburgh, September 1986.
- [Chapman, 1988] David Chapman. How to do research at the MIT AI lab. Technical Report AI Working Paper 316, MIT, October 1988.
- [Division, 1986a] UC Berkeley Graduate Division. Finding money for dissertation research/writing. *The Graduate*, II(3), Fall 1986.
- [Division, 1986b] UC Berkeley Graduate Division. Studying for the qualifying exam. *The Graduate*, II(3), Fall 1986.
- [Division, 1986c] UC Berkeley Graduate Division. Writing your thesis. *The Graduate*, II(1), Spring 1986.
- [Division, 1987a] UC Berkeley Graduate Division. Interviewing for a faculty position. *The Graduate*, III(2), Fall 1987.
- [Division, 1987b] UC Berkeley Graduate Division. The making of a successful proposal. *The Graduate*, III(1), Spring 1987.
- [Division, 1988] UC Berkeley Graduate Division. Choosing your thesis or dissertation topic. *The Graduate*, IV(2), Fall 1988.
- [Division, 1989] UC Berkeley Graduate Division. Beating the isolation blues. *The Graduate*, V(1), Spring 1989.
- [Hall and Sandler, a] Roberta M. Hall and Bernice R. Sandler. Academic mentoring for women students and faculty: A new look at an old way to get ahead.
- [Hall and Sandler, b] Roberta M. Hall and Bernice R. Sandler. The classroom climate: A chilly one for women?
- [Hall and Sandler, c] Roberta M. Hall and Bernice R. Sandler. Out of the classroom: A chilly campus climate for women?
- [Lefferts, 1978] Robert Lefferts. *Getting a Grant: How to Write Successful Grant Proposals*. 1978.
- [Leveson, 1989] Nancy Leveson. Women in computer science: A report for the NSF CISE cross-disciplinary activities advisory committee, December 1989.
- [McKay, 1988] Nellie Y. McKay. Minority faculty in [mainstream white] academia, 1988. Chapter 5.
- [NSF, 1989] NSF. An NSF study and report about women in computing research. *Computing Research News*, Summer 1989.
- [Sandler and Hall, 1986] Bernice R. Sandler and Roberta M. Hall. The campus climate revisited: Chilly for women faculty, administrators, and graduate students, October 1986.

- [Shriver, 1990] Bruce D. Shriver. The benefits of quality refereeing. *COMPUTER*, pages 10–16, April 1990. Also includes *COMPUTER*'s guidelines for referees.
- [Smith, 1990] Alan Jay Smith. The task of the referee. *COMPUTER*, pages 65–71, April 1990.
- [Somerville, 1982] Bill Somerville. Where proposals fail: A foundation executive's basic list of what to do and not do when requesting funding. *The Grantsmanship Center News*, Jan/Feb 1982.
- [Spertus, 1992] Ellen Spertus. Why are there so few female computer scientists?, 1992. Expected to become an MIT AI Lab Technical Report.
- [Strok, 1992] Dale Strok. Women in AI. *IEEE Expert*, 7(4):7–21, August 1992.
- [Toth, 1988] Emily Toth. Women in academia. In *The Academics' Handbook*. Duke University Press, 1988. Chapter 4.
- [White, 1975] Virginia White. *Grants: How to Find Out About Them and What to Do Next*. Plenum Press, 1975.
- [White, 1983] Virginia White. *Grant Proposals That Succeeded*. Plenum Press, 1983.